Preparation contextuality powers parity-oblivious multiplexing

Robert W. Spekkens,[1] D. H. Buzacott,[2,3] A. J. Keehn,[2,3] Ben Toner,[4] and G. J. Pryde[2,3]

[1]*DAMTP, University of Cambridge, Cambridge, United Kingdom CB3 0WA*
[2]*Centre for Quantum Computer Technology, Griffith University, Brisbane 4111, Australia*
[3]*Centre for Quantum Dynamics, Griffith University, Brisbane 4111, Australia*
[4]*Centrum voor Wiskunde en Informatica, Kruislaan 413, 1098 SJ Amsterdam, The Netherlands*

In a noncontextual hidden variable model of quantum theory, hidden variables determine the outcomes of every measurement in a manner that is independent of how the measurement is implemented. Using a generalization of this notion to arbitrary operational theories and to preparation procedures, we demonstrate that a particular two-party information-processing task, "parity-oblivious multiplexing," is powered by contextuality in the sense that there is a limit to how well any theory described by a noncontextual hidden variable model can perform. This bound constitutes a "noncontextuality inequality" that is violated by quantum theory. We report an experimental violation of this inequality in good agreement with the quantum predictions. The experimental results also provide the first demonstration of 2-to-1 and 3-to-1 quantum random access codes.

PACS numbers: 03.65.Ta, 03.67.-a, 42.50.Dv, 42.50.Ex, 42.50.Xa

arXiv:0805.1463v2 [quant-ph] 8 Jan 2009

The Bell-Kochen-Specker theorem [1] shows that the predictions of quantum theory are inconsistent with a hidden variable model having the following feature: if A, B and C are Hermitian operators such that A and B commute, A and C commute, but B and C do not commute, then the value predicted to occur in a measurement of A does not depend on whether B or C was measured simultaneously. This feature is called "noncontextuality." Significantly, it is only well-defined for models of quantum theory (and then only for projective measurements and deterministic models) [2]. By contrast, Bell's definition of a *local* model applies to any theory that can be described operationally [3]. Consequently, whereas one can test whether or not experimental statistics are consistent with a local model (by testing whether or not they satisfy Bell inequalities), there is no way to test whether or not experimental statistics are consistent with a noncontextual model (and no way of defining associated "noncontextuality inequalities") *unless* one generalizes the traditional notion of noncontextuality in such a way that it makes no reference to the quantum formalism. Suggestions for such a formulation have been made by several authors [4]. A particularly natural generalization (and slight modification) which applies to all models (deterministic or not) of any operational theory has been proposed in Ref. [2]. We here derive a noncontextuality (NC) inequality based on this notion.

Because information-theoretic tasks can be characterized entirely in terms of experimental statistics, one can explore whether theories that violate NC inequalities may provide information-theoretic advantages over theories that satisfy these inequalities. We prove that this is indeed the case for a task which we call *parity-oblivious multiplexing*, a kind of two-party secure computation. (The notion that contextuality might yield an advantage for multiplexing tasks was first put forward by Galvão [5].) The NC inequality we derive provides a bound on the probability of success in this task and we demonstrate a quantum protocol for parity-oblivious multiplexing for which the probability of success exceeds the noncontextual bound.

Finally, we report an experimental implementation of this protocol that achieves a probability of success in good agreement with the quantum result and in violation of the NC inequality.

Operational theories and noncontextual models. In an operational theory, the primitives of description are preparations and measurements, specified as instructions for what to do in the laboratory. The theory simply provides an algorithm for calculating the probability $p(k|P,M)$ of an outcome k of measurement M given a preparation P. As an example, in quantum theory, every preparation P is represented by a density operator ρ_P, every measurement M is represented by a positive operator valued measure $\{E_{M,k}\}$, and the probability of outcome k is given by $p(k|P,M) = \mathrm{Tr}\left(\rho_P E_{M,k}\right)$.

In a hidden variable model of an operational theory, a preparation procedure is assumed to prepare a system with certain properties and a measurement procedure is assumed to reveal something about those properties. The set of all variables describing the system is denoted λ. It is presumed that for every preparation P, there is a probability distribution $p(\lambda|P)$ such that implementing P causes the system to be prepared in physical state λ with probability $p(\lambda|P)$. Similarly, it is presumed that for every measurement M, there is a distribution $p(k|\lambda,M)$ such that implementing M on a system described by λ yields outcome k with probability $p(k|\lambda,M)$. For the hidden variable model to reproduce the predictions of the operational theory, it must satisfy $p(k|P,M) = \int \mathrm{d}\lambda p(k|\lambda,M)p(\lambda|P)$.

A hidden variable model is *preparation noncontextual*

if the following implication holds

$$\forall M : p(k|P,M) = p(k|P',M) \rightarrow \ p(\lambda|P) = p(\lambda|P'), \tag{1}$$

that is, if two preparations yield the same statistics for all possible measurements then they are represented equivalently in the hidden variable model. Similarly, *measurement noncontextuality* is the condition that

$$\forall P : p(k|P,M) = p(k|P,M') \rightarrow \ p(k|\lambda,M) = p(k|\lambda,M'), \tag{2}$$

that is, if two measurements have the same statistics for all possible preparations then they are represented equivalently in the model. More details can be found in Ref. [2]. An NC inequality is any inequality on experimental statistics that follows from the assumption that there exists a hidden variable model that is preparation and measurement noncontextual. It is of the form $f(p(k|P_1,M_1),p(j|P_2,M_2),...) \leq C$ for some function f and constant C.

Parity-oblivious multiplexing. Suppose that Alice and Bob wish to perform the following information-processing task, which we call *n-bit parity-oblivious multiplexing*. Alice has as input an n-bit string x chosen uniformly at random from $\{0,1\}^n$. Bob has as input an integer y chosen uniformly at random from $\{1,\ldots,n\}$ and must output the bit $b = x_y$, that is, the yth bit of Alice's input. Alice can send a system to Bob encoding information about her input, however there is a cryptographic constraint: no information about any parity of x can be transmitted to Bob. More specifically, letting $s \in \text{Par}$ where $\text{Par} \equiv \{r | r \in \{0,1\}^n, \sum_i r_i \geq 2\}$ is the set of n-bit strings with at least two bits that are 1, no information about $x \cdot s = \bigoplus_i x_i s_i$ (termed the *s-parity*) for any such s can be transmitted to Bob (here $\oplus$ denotes sum modulo 2). This task is similar to an n-to-1 quantum random access code [5, 6, 7, 8] except that it has a constraint of parity-obliviousness rather than a constraint on the potential information-carrying capacity of the system used.

Lemma 1. *Classically, the optimal probability of success in n-bit parity-oblivious multiplexing satisfies* $p(b = x_y) \leq (n+1)/2n$.

Proof. (For details, see Appendix A.) The only classical encodings of x that reveal no information about any parity (while encoding *some* information about x) are those that encode only a single bit x_i for some i. Given that y is uniformly distributed, it makes no difference which bit it is. Therefore, we may assume that Alice and Bob agree that Alice will always encode the first bit, x_1. If $y = 1$, which occurs with probability $1/n$, then Bob can output $b = x_y$ and win. With probability $(n-1)/n$, we have $y \neq 1$ and in this case Bob can at best guess the value of x_y and wins with probability $1/2$. ■

What is the most general protocol that can be implemented in an arbitrary operational theory? For each input string x, Alice implements a preparation procedure P_x, and for each integer y, Bob implements a binary-outcome measurement M_y, and reports the outcome b as his output. The probability of winning is

$$p(b = x_y) = \frac{1}{2^n n} \sum_{y \in \{1,\ldots,n\}} \sum_{x \in \{0,1\}^n} p(b = x_y | P_x, M_y) \tag{3}$$

where $1/2^n n$ is the prior probability for a particular x and y. The parity-oblivious constraint requires that for every s-parity, there is no outcome of any measurement for which posterior probabilities for s-parity 0 and s-parity 1 are different, that is,

$$\forall s \forall M \forall k : \sum_{x|x\cdot s=0} p(P_x|k,M) = \sum_{x|x\cdot s=1} p(P_x|k,M). \tag{4}$$

Noncontextuality inequality. The main theoretical result of this letter is the following theorem.

Theorem 2. *In an operational theory that admits a preparation noncontextual hidden variable model, the optimal probability of success in n-bit parity-oblivious multiplexing satisfies* $p(b = x_y) \leq (n+1)/2n$.

Proof. Define $P_{s,b}$ to be the procedure obtained by choosing uniformly at random an x such that $x \cdot s = b$ and implementing P_x. Clearly, for any measurement M, the probability of outcome k given preparation $P_{s,b}$ is simply

$$p(k|P_{s,b},M) = \frac{1}{2^{n-1}} \sum_{x|x\cdot s=b} p(k|P_x,M). \tag{5}$$

Similarly, the probability of hidden variable λ given an implementation of $P_{s,b}$ is simply

$$p(\lambda|P_{s,b}) = \frac{1}{2^{n-1}} \sum_{x|x\cdot s=b} p(\lambda|P_x). \tag{6}$$

Now note that one can re-express the parity-oblivious condition, Eq. (4), as $\forall s \forall M : \sum_{x|x\cdot s=0} p(k|P_x,M) = \sum_{x|x\cdot s=1} p(k|P_x,M)$ (it follows from Bayes' rule and the uniformity of the prior over x). Combining this with Eq. (5), we infer that $\forall s \forall M : p(k|P_{s,0},M) = p(k|P_{s,1},M)$ which is simply the statement that mixed preparations corresponding to opposite s-parities are indistinguishable by any measurement. But together with the assumption that the hidden variable model is preparation noncontextual, Eq. (1), this implies that $\forall s : p(\lambda|P_{s,0}) = p(\lambda|P_{s,1})$, which states that mixed preparations corresponding to opposite s-parities are also indistinguishable at the hidden variable level. Using Eq. (6) and Bayes' rule again, we obtain

$$\forall s : \sum_{x|x\cdot s=0} p(P_x|\lambda) = \sum_{x|x\cdot s=1} p(P_x|\lambda). \tag{7}$$

Therefore, even if one knew λ, the posterior probabilities for s-parity 0 and s-parity 1 would be the same, that is,

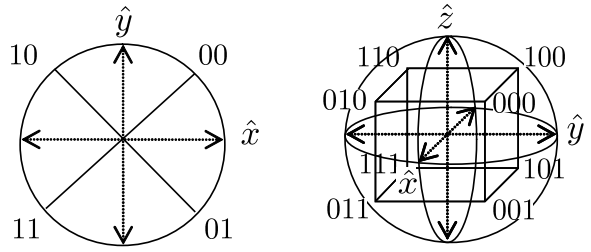

FIG. 1: Bloch representation of states and measurements in quantum 2-bit and 3-bit parity-oblivious multiplexing.

one would know nothing about any s-parity of x. The argument so far can be summarized as follows: for preparation noncontextual models, parity-obliviousness at the operational level implies parity-obliviousness at the level of the hidden variables.

The hidden state λ provides a classical encoding of x. But, as just shown, it is one that cannot contain information about any s-parity. We recall from lemma 1 that such encodings have information about at most one bit, x_i, of x. Consequently, even if Bob could determine λ perfectly, he and Alice could at best achieve the optimal probability of success achievable in a classical protocol (specified in lemma 1), while if Bob is limited in his ability to determine λ (as will be the case in general in a hidden variable model), they will do worse. ■

Quantum case. We now consider how well one can achieve parity-oblivious multiplexing in quantum theory. The following is a protocol for the 2-bit case that uses a single qubit as the quantum message. Alice encodes her 2 bits into the four pure quantum states with Bloch vectors $(\pm\frac{1}{\sqrt{2}}, \pm\frac{1}{\sqrt{2}})$ equally distributed on an equatorial plane of the Bloch sphere, as indicated in Fig. 1 (recall that a density operator ρ is related to its Bloch vector $\vec{r}$ by $\rho = \frac{1}{2}(I + \vec{r}\cdot\vec{\sigma})$, where $\vec{\sigma}$ is the vector of Pauli matrices). Bob measures along the $\hat{x}$ axis if he wishes to learn the first bit, and along the $\hat{y}$ axis if he wishes to learn the second. He guesses the bit value 0 upon obtaining the positive outcome. In all cases, the guessed value is correct with probability $\cos^2(\pi/8) \simeq 0.853553$. Meanwhile, no information about the parity can be obtained by any quantum measurement given that the parity 0 and parity 1 mixtures are represented by the same density operator, $\frac{1}{2}\rho_{00} + \frac{1}{2}\rho_{11} = \frac{1}{2}\rho_{01} + \frac{1}{2}\rho_{10} = I/2$. We have a violation of the NC inequality of Thm. 2 because for $n = 2$, the upper bound on the probability of success is 3/4. By exploiting a connection with the Clauser-Horne-Shimony-Holt inequality [9], one can show that this protocol yields the maximum possible quantum violation of the NC inequality.

A protocol for 3-bit parity-oblivious multiplexing using a single qubit proceeds as follows. Alice encodes her three bits into a set of eight pure quantum states associated with Bloch vectors $(\pm\frac{1}{\sqrt{3}}, \pm\frac{1}{\sqrt{3}}, \pm\frac{1}{\sqrt{3}})$ forming a cube inside the Bloch sphere (see Fig. 1). Bob measures along the $\hat{x}$, $\hat{y}$ or $\hat{z}$ axes to obtain the first, second or third bits. In all cases, the guessed value is correct with probability $\frac{1}{2}(1+\frac{1}{\sqrt{3}}) \simeq 0.788675$. The mixture of the four states corresponding to $x_1 \oplus x_2 = 0$ (i.e. s-parity 0 for $s = (1, 1, 0)$) is identical to the mixture of the four states corresponding to $x_1 \oplus x_2 = 1$ and is equal to $I/2$. Similarly for the two mixtures associated with each of the other three parities, $x_1 \oplus x_3$ ($s = (1, 0, 1)$), $x_2 \oplus x_3$ ($s = (0, 1, 1)$), and $x_1 \oplus x_2 \oplus x_3$ ($s = (1, 1, 1)$). The protocol is therefore parity-oblivious for all s-parities. Again we have a violation of the NC inequality because for $n = 3$ the upper bound on the probability of success is 2/3. It is an open question whether 0.788675 is the maximum possible quantum violation.

The 2-bit protocol was originally presented as a 2-to-1 quantum random access code by Wiesner [6] and rediscovered in Ref. [7], while the 3-bit protocol was presented in Ref. [8] as an instance of a 3-to-1 quantum random access code (the original idea is attributed to Chuang in Ref. [7]).

Experimental results. We experimentally demonstrate better-than-classical performance for 2-bit and 3-bit parity-oblivious multiplexing by implementing the quantum protocols using polarization qubits. Photon pairs from downconversion are coupled into single mode optical fibers. One photon acts as a trigger, while the other is used in the experiment. Alice's state preparation consists of a fiber polarization controller, and a polarizing beam displacer, rotated to the input state angle, used to ensure high-purity linearly polarized states for the 2-bit protocol. An additional quarter wave plate is used to prepare elliptically-polarized states for the 3-bit protocol. Bob's measurement consists of a polarizing beam displacer mounted in a computer-controlled rotation mount, followed by a single photon counting module. For our demonstration, a detector is placed at only a single output port of the beam displacer and the probability of each outcome is calculated from the relative number of counts for a given beam displacer angle and the one orthogonal to it. (Further details of the experimental set-up, including a figure, are provided in Appendix B.) Adjustment of the beam displacer and quarter wave plate angles allows measurement of the horizontal/vertical basis, the diagonal/anti-diagonal basis and the right/left-circular basis. Valid measurement events are heralded by a coincidence count between the directly detected photon and the experiment photon. These experimental procedures for a given x and y define the preparation P_x and the measurement M_y respectively.

We obtained probabilities $p(k = x_y|P_x, M_y)$ by accumulating statistics over approximately 3.5×10^7 coincidence counts for each x and y in the 2-bit scheme and 2.4×10^7 in the 3-bit scheme. Using Eq. (3), we cal-

culated the 2-bit and 3-bit probabilities of success to be $p(b = x_y) = 0.851929 \pm 0.000030$ and $p(b = x_y) = 0.786476 \pm 0.000017$ respectively. The errors were determined from the Poissonian counting statistics of the parametric source and the small repeatability error in the wave plate settings, using standard error analysis techniques. These probabilities of success violate the NC inequality of Thm. 2 with a high degree of confidence: 3410 and 6922 standard deviations respectively. They are also close to the predicted quantum values of 0.853553 and 0.788675, achieving a violation that is 98.4% and 98.2% respectively of the gap between the NC bound and the quantum value.

Just as Bell inequality violations are only surprising given the absence of signalling between the two wings of the experiment, the NC inequality violations are only surprising given the parity-oblivious property. However, whereas one can establish the absence of signalling by confirming that the two wings are space-like separated, one must directly test for transmission of information about the parity in our experiment. A consideration of how this is to be accomplished highlights two shortcomings in the operational definition of preparation noncontextuality of Eq. (1): in practice one can never implement *all* measurements and one never finds truly *identical* statistics. The first issue may be addressed by relying on previous experimental evidence for the existence of a tomographically complete set of measurements – one from which the statistics of any other measurement can be calculated – and testing indistinguishability relative to this set alone, as we shall do here. The second issue may be addressed by presuming a kind of continuity: closeness of experimental statistics implies closeness of the representations in the model [2] (this parallels the problem of dealing with imperfect alignment in traditional proofs of contextuality [12], where continuity also provides an answer [4, 13]). In the present work, we simply demonstrate that the experimental statistics are *close* to parity-oblivious while yielding a large violation of the noncontextuality inequalities, and leave a more detailed analysis for future work.

We quantify the obliviousness of our experimental protocol for a particular s-parity by the maximum probability that Bob can correctly estimate this parity in a variation over all measurements. One can estimate this by implementing a tomographically complete set of measurements, then reconstructing the states ρ_0 and ρ_1 associated with s-parity 0 and s-parity 1, and finally making use of the fact that the maximum probability of discriminating these states is $\frac{1}{2} + \frac{1}{4}\mathrm{Tr}|\rho_0 - \rho_1|$. Among all s-parities, we calculate the largest such probability to be 0.5020 ± 0.0002. This calculation is not sufficient, however, because it neglects an imperfection in the experiment that also contributes to leakage of information about the parity, namely, that there is a small probability of more than one photon being sent to the experiment. By our characterization of the source, we estimate the probability of two photons to be 0.007 ± 0.003 relative to the single photon generation probability. If two photons pass through the polarizers in the ideal protocol, the maximum probability of correctly estimating the parity can be quite far from 1/2: it is 3/4 in the case of the 2-bit scheme and 2/3 for three of the four s-parities in the 3-bit scheme. However, the fact that this possibility occurs with low probability implies that the two-photon contribution to the probability of correct estimation is comparable to the one-photon contribution. (Contributions from three or more photons are negligible in comparison). The weighted average of these contributions is easily calculated and the largest, among all s-parities, is found to be 0.504 ± 0.002. The fact that this is within one percent of 1/2 demonstrates that our experimental protocols are indeed close to parity-oblivious.

Given that the quantum protocols described herein are also 2-to-1 and 3-to-1 random access codes, our results constitute the first experimental demonstration of a quantum advantage for these tasks as well.

Finally, it is worth noting that every Bell inequality is a special case of an NC inequality where all assumptions of noncontextuality are justified by locality [2]. Consequently, every experimental violation of a Bell inequality demonstrates the impossibility of a noncontextual hidden variable model. Indeed, this is *all* that can be demonstrated by those that fail to seal the locality loophole [10, 11]. Nonetheless, a dedicated experiment of the sort we have described here can achieve a large violation with high confidence at a smaller cost of experimental effort.

Acknowledgements. R.W.S. thanks M. Leifer and J. Barrett for helpful discussions. This work has been supported by the Australian Research Council, an IARPA-funded US Army Research Office contract, NWO VICI project 639-023-302, the Dutch BSIK/BRICKS project, the EUs FP6-FET Integrated Projects SCALA (CT-015714) and QAP (CT-015848), and the Royal Society.

APPENDIX A: OPTIMAL CLASSICAL PROTOCOL FOR N-BIT PARITY-OBLIVIOUS MULTIPLEXING

We here provide a more detailed proof of lemma 1. First, note that by the assumption of parity-obliviousness, the classical message m sent from Alice to Bob must satisfy

$$\forall s : \sum_{x|x\cdot s=0} p(P_x|m) = \sum_{x|x\cdot s=1} p(P_x|m) \qquad (8)$$

By Bayes' theorem and the fact that the distribution over inputs x is uniform, we can rewrite this as a constraint

on $p(m|P_x)$, namely,

$$\forall s : \sum_{x|x\cdot s=0} p(m|P_x) = \sum_{x|x\cdot s=1} p(m|P_x). \tag{9}$$

As we will demonstrate (at the end of this section), this implies that $p(m|P_x)$ has the form

$$p(m|P_x) = p(0)p_0(m) + \sum_{i=1}^{n} p(i)\left[p_{i,0}(m)\delta_{x_i,0} + p_{i,1}(m)\delta_{x_i,1}\right], \tag{10}$$

where $p(i)$ is a normalized probability distribution on $\{0,\ldots,n\}$, the functions $p_0(m), p_{i,0}(m)$ and $p_{i,1}(m)$ are normalized probability distributions over m, and where $\delta_{a,b}$ is the Kronecker delta function (equal to 1 if $a=b$ and 0 otherwise).

It follows that any classical parity-oblivious multiplexing protocol can be interpreted as follows: Alice generates an integer $i \in \{0,\ldots,n\}$ from the distribution $p(i)$. Upon obtaining $i=0$, she sends a message m chosen from the distribution $p_0(m)$ (independent of the value of x). Upon obtaining $i \in \{1,\ldots,n\}$, she sends a message m chosen from one of two distributions, depending on the value of the ith bit of x : the distribution is $p_{i,0}(m)$ if $x_i = 0$ and $p_{i,1}(m)$ if $x_i = 1$.

We now determine the choice of these distributions that leads to a maximum probability of winning. First note that if $i=0$, Bob gets no information about x. This is clearly not optimal, so we may set $p(0)=0$. Next note that the amount that Bob learns about x_i depends on his ability to distinguish $p_{i,0}(m)$ from $p_{i,1}(m)$. To optimize the amount that Bob can learn, $p_{i,0}(m)$ and $p_{i,1}(m)$ must be chosen to be perfectly distinguishable. This is only possible if they are completely non-overlapping, that is, if $p_{i,0}(m)p_{i,1}(m)=0$.

In an optimal decoding, Bob simply determines whether m is in the support of $p_{y,0}(m)$ or of $p_{y,1}(m)$ and outputs $b=0$ or 1 accordingly. This is optimal for the following reason. The message m only contains information about x_y if Alice happened to generate an i that coincides with y and in this case Bob will output $b=x_y$ with probability 1. When i does not coincide with y, Bob gets no information about x_y from m, so it is irrelevant what he outputs; given that x_y is equally likely to be 0 or 1, his probability of having generated the correct output will be 1/2.

Finally, given that y is chosen uniformly at random, the probability of i coinciding with y is $1/n$, so that the overall probability of a correct output is $\frac{1}{n}(1) + \left(1-\frac{1}{n}\right)\left(\frac{1}{2}\right) = (n+1)/2n$.

It is worth noting that there are many natural schemes that achieve the optimum:

- If $p(i) = \delta_{i,j}$ for some particular $j \in \{1,\ldots,n\}$, and $p_{j,0}(m)p_{j,1}(m)=0$, then Alice has simply encoded the value of x_j in her message.
- For $p(i)$ an arbitrary distribution over $\{1,\ldots,n\}$, if $p_{i,0}(m)p_{i,1}(m)=0$ for all i, then Alice has simply chosen a value $i \in \{1,\ldots,n\}$ according to this distribution and encoded x_i in her message.
- For $p(i)$ an arbitrary distribution over $\{1,\ldots,n\}$, if $p_{i,b}(m)p_{i',b'}(m)=0$ when either $b \neq b'$ or $i \neq i'$, then Alice has encoded both i and x_i in her message.

It remains to prove our claim that the parity-oblivious constraint, Eq. (9), implies the decomposition of $p(m|P_x)$ described in Eq. (10). We do this using Fourier analysis over $\mathbb{Z}_2^n$. Let $r \in \{0,1\}^n$. Define functions $\chi_r : \{0,1\}^n \to [-1,1]$ where

$$\chi_r(x) := (-1)^{x\cdot r}.$$

These form an orthonormal set because

$$\sum_x \chi_r(x)\chi_{r'}(x) = \sum_{x\in\{0,1\}^n} (-1)^{x\cdot(r\oplus r')} = 2^n\delta_{r,r'}.$$

Moreover, noting that the dimensionality of the space of functions on $\{0,1\}^n$ is 2^n (a parameter for every input string) and that there are 2^n values of r, we see that the χ_r form an orthonormal basis of the function space. It follows that we can write $p(m|P_x)$ in the Fourier series

$$p(m|P_x) = \sum_r \hat{p}(m,r)\chi_r(x).$$

We infer that

$$2^n\hat{p}(m,r) = \sum_x \chi_r(x)p(m|P_x) = \sum_{x|x\cdot r=0} p(m|P_x) - \sum_{x|x\cdot r=1} p(m|P_x).$$

Combining this with the parity-obliviousness condition, Eq. (9), one obtains

$$\forall s \in \mathrm{Par} : \hat{p}(m,s) = 0.$$

Consequently, the only strings r for which $\hat{p}(m,r) \neq 0$ are those with Hamming weight 0 or 1. Denoting the Fourier coefficients of the all zero string by $\hat{p}_0(m)$ and that of the string with a single 1 at position i by $\hat{p}_i(m)$, we have

$$p(m|P_x) = \hat{p}_0(m) + \sum_{i=1}^{n} \hat{p}_i(m)(-1)^{x_i}.$$

Because $(-1)^{x_i} = \delta_{x_i,0} - \delta_{x_i,1}$ and $1 = \delta_{x_i,0} + \delta_{x_i,1}$, we can write

$$p(m|P_x) = a_0(m) + \sum_{i=1}^{n} \left[a_{i,0}(m)\delta_{x_i,0} + a_{i,1}(m)\delta_{x_i,1}\right] \tag{11}$$

where we have defined nonnegative coefficients

$$a_{i,0}(m) = 2\hat{p}_i(m),\ a_{i,1}(m) = 0 \qquad \text{if } \mathrm{sgn}(\hat{p}_i(m)) \geq 0,$$
$$a_{i,0}(m) = 0,\ a_{i,1}(m) = -2\hat{p}_i(m) \qquad \text{if } \mathrm{sgn}(\hat{p}_i(m)) < 0;$$

and we have implicitly defined a constant $a_0(m)$, which we presently show is also nonnegative. To do this, we define an n-bit string $z(m)$ that encodes the signs of the Fourier coefficients. Specifically, $z(m)$ is defined by

$$z_i(m) \equiv \begin{cases} 1 & \text{if } \mathrm{sgn}(\hat{p}_i(m)) \geq 0, \\ 0 & \text{if } \mathrm{sgn}(\hat{p}_i(m)) < 0. \end{cases}$$

It follows from this definition that

$$a_{i,0}(m)\delta_{z_i(m),0} + a_{i,1}(m)\delta_{z_i(m),1} = 0$$

for all i, and consequently that

$$p(m|P_{z(m)}) = a_0(m),$$

which establishes that $a_0(m) \geq 0$.

Finally, we show that Eq. (11) can be put into the form of Eq. (10). By the normalization of the distribution $p(m|P_x)$, we have

$$1 = \sum_m p(m|P_x) = \sum_m a_0(m) + \sum_{i=1}^{n} \sum_m a_{i,x_i}(m),$$

for all x. Defining $A_0 = \sum_m a_0(m)$ and $A_{i,x_i} = \sum_m a_{i,x_i}(m)$, we have

$$A_0 + \sum_{i=1}^{n} A_{i,x_i} = 1,$$

for all x, which implies that $\sum_{i=1}^{n} A_{i,x_i}$ is independent of x and in particular of x_i. We deduce that

$$A_{i,0} = A_{i,1}$$

for all i. Eq. (10) now follows from Eq. (11) by identifying

$$p(0) = A_0$$
$$p(i) = A_{i,0} = A_{i,1}$$

and

$$p_0(m) = a_0(m)/p(0)$$
$$p_{i,b}(m) = a_{i,b}(m)/p(i).$$

if $p(0), p(i) \neq 0$.

APPENDIX B: EXPERIMENTAL DETAILS

A schematic of the experimental set-up is provided in Fig. 2. We used type-I downconversion in bismuth borate (BiBO) to generate pairs of 820 nm, horizontally polarized single photons from a 410 nm, 60 mW continuous-wave diode laser. A 10 nm FWHM interference filter is used to reject background light. In the experiment, we obtained coincidence rates (2.5 ns window) of approximately 23100 pairs/s in the 2-bit scheme and 15200 pairs/s in the 3-bit scheme.

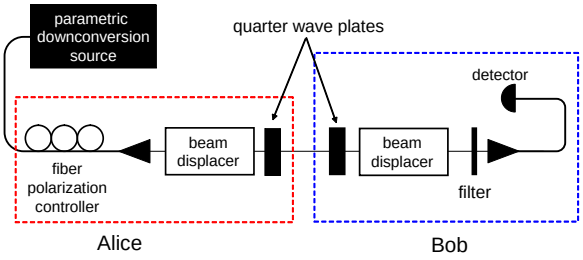

FIG. 2: Experimental set-up. The parametric downconversion source provides single photons to the experiment. Detection events in the experiment are counted in coincidence with the downconversion trigger photon (not shown).

Although we chose to implement the experiment with a heralded mode of a downconversion source, similar results could also have been obtained with weak coherent states (using the same measurements) [15]. In both cases, one must postselect on not finding the vacuum (implying, incidentally, that the detector loophole is not sealed [14]), and in both cases there is a small amplitude for more than one photon and hence a small amount of leaked parity information. Our choice was motivated by differences in ideal performance – it is only for the downconversion scheme that the leakage of parity information can be eliminated in principle (through the use of true single photons heralded by efficient number-resolving detectors). Nonetheless, this ideal has not yet been realized.

[1] J. S. Bell, Rev. Mod. Phys. **38**, 447 (1966); S. Kochen and E. P. Specker, J. Math. Mech. **17**, 59 (1967).
[2] R. W. Spekkens, Phys. Rev. A **71**, 052108 (2005).
[3] J. S. Bell, Physics **1**, 195 (1964).
[4] A. Cabello and G. Garcia-Alcaine, Phys. Rev. Lett. **80,** 1797 (1998); C. Simon, C. Brukner and A. Zeilinger, Phys. Rev. Lett. **86,** 4427 (2001); J.-A. Larsson, Europhys. Lett. **58**, 799 (2002).
[5] E. F. Galvão, Ph.D. thesis, arXiv:quant-ph/0212124v1.
[6] S. Wiesner, Sigact News **15**, 78 (1983).
[7] A. Ambainis, A. Nayak, A. Ta-Shma, and U. Vazirani, in Proceedings of the 31st Annual ACM Symposium on the Theory of Computing (ACM Press, New York, 1999).
[8] M. Hayashi *et al.*, New Journal of Physics **8**, 129 (2006).
[9] J. F. Clauser, M.A. Horne, A. Shimony and R. A. Holt, Phys. Rev. Lett. **23**, 880 (1969).
[10] M. A. Rowe *et al.*, Nature **409**, 791 (2001).

[11] Y. Hasegawa *et al.*, Nature **425**, 45 (2003).
[12] D. A. Meyer, Phys. Rev. Lett. **83** 3751 (1999).
[13] N. D. Mermin, arXiv:quant-ph/9912081.
[14] P. Pearle, Phys. Rev. D, **2**, 1418 (1970).
[15] The Wigner representation would not provide a classical statistical model of the experiment because the representation of the measurements would be nonpositive.